



Education

Desktop, Nearline & Enterprise HDDs

What's the difference?

Willis Whittington, Seagate Technology

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced without modification
 - ◆ The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.

Desktop, Nearline & Enterprise Disk Drives What's the difference?

For the past twenty five years the storage marketplace has been divided into two major categories namely “Desktop” and “Enterprise”. Recently, a third player variously known as “Nearline”, “Reference” or “Business Critical” has evolved to provide a low cost, high capacity storage solution for Enterprise data that no longer needs to exist in a high availability transactional processing environment but must maintain 24 x 7 availability as a reference or backup resource.

Each of these classes of drives requires a unique and specific set of attributes to fulfill its role. This presentation will explore these differences and explain why you need to use

the right drive for the right application.

- Basic Comparisons
- SAS & SATA Compatibility
- The Advantages of Nearline SAS
- Rotational Vibration
- Data Error Rate
- Error Correction Capability
- Data Integrity
- Performance
- Annualized Failure Rate

~ Q & A along the way ~

Basic Comparisons

Comparison Table DT / NL / MC*


Key: Good Better Best

Metric	Desktop	SATA Nearline (NL)	SAS Nearline (NL)	Enterprise MC*
Capacity (GB)	1,000	1,000	1,000	450
Cost	Low	Mid	Mid	High
Power Consumption	1x	1.2x	1.2x+	1.5x
MTBF (Hrs)	600,000	1,200,000	1,200,000	1,600,000
Duty Cycle	Low (<10%)	Low/Medium (<20%)	Low/Medium (<20%)	High (100%)
Data Integrity	Parity (?)	EDC + (ECC?)	EDC/ECC + Proprietary Data Integrity Protection	EDC/ECC + Proprietary Data Integrity Protection
Unrec Error Rate	10 ⁻¹⁴	10 ⁻¹⁵	10 ⁻¹⁵	10 ⁻¹⁶
RV Radians/sec ²	6	12.5	12.5	>21
Error Recovery	SATA	SATA + Time Control	Full SCSI	Full SCSI
Firmware/Features	Standard SATA	SATA + Selected Nearline Features	SCSI + Adv. Features (Enabled by Dual CPU)	SCSI + Adv. Features (Enabled by Dual CPU)
Power On Hrs/Year	2400	8,760 (Low Duty cycle)	8,760 (Low Duty cycle)	8,760
Multi Initiator	No	No	16 Hosts & Dual Port	16 Hosts & Dual Port
Performance	1x	1x	1x+	1.4x / 2.5x (Seq / Rand)
T10 Data Protection	No	No	Yes	Yes
Scalability	Low	Low	High + Dual Port	High + Dual Port

*Mission Critical

Desktop, Nearline & Enterprise HDDs – What's the difference?
© 2008 Storage Networking Industry Association. All Rights Reserved.


Anatomy of an Enterprise Drive



Motor

Higher rpm than NL or DT
Tighter specifications
Less runout
More expensive

Anatomy of an Enterprise Drive



Discs


Four platter design

Smaller diameter than NL/DT


Full media certification

Fully characterized

Variable sector format



Anatomy of an Enterprise Drive



Head Stack


Eight head design

Low mass, high rigidity


Voice coil designed for

- optimal performance
- 100% duty cycle

Higher cost design




Anatomy of an Enterprise Drive



Environmental Control Module

- Humidity Control
- Chemical Absorbtion
- Multi-point filtration
- Windage Design

Anatomy of an Enterprise Drive



Misc Mechanical

- Powerful Voice Coil Motor
- Stiffer Covers
- Air Control Devices
- Faster Seeks
- High Servo Sample Rate
- Low RV susceptibility

Anatomy of an Enterprise Drive




Electronics


- Dual processors
- Multi host queuing
- Dual port
- Twice the memory of NL/DT
- High rpm control
- Command scheduling
- Superior error protection
- Superior error correction
- Smart servo algorithms
- Perform. optimization
- Data integrity checks
- Sequential h/w assist

SAS & SATA Compatibility


I/O Connectors for SAS & SATA



Current Limited



Steady State



For SAS, the key-way is filled in and its flip side is used for the 2nd Port

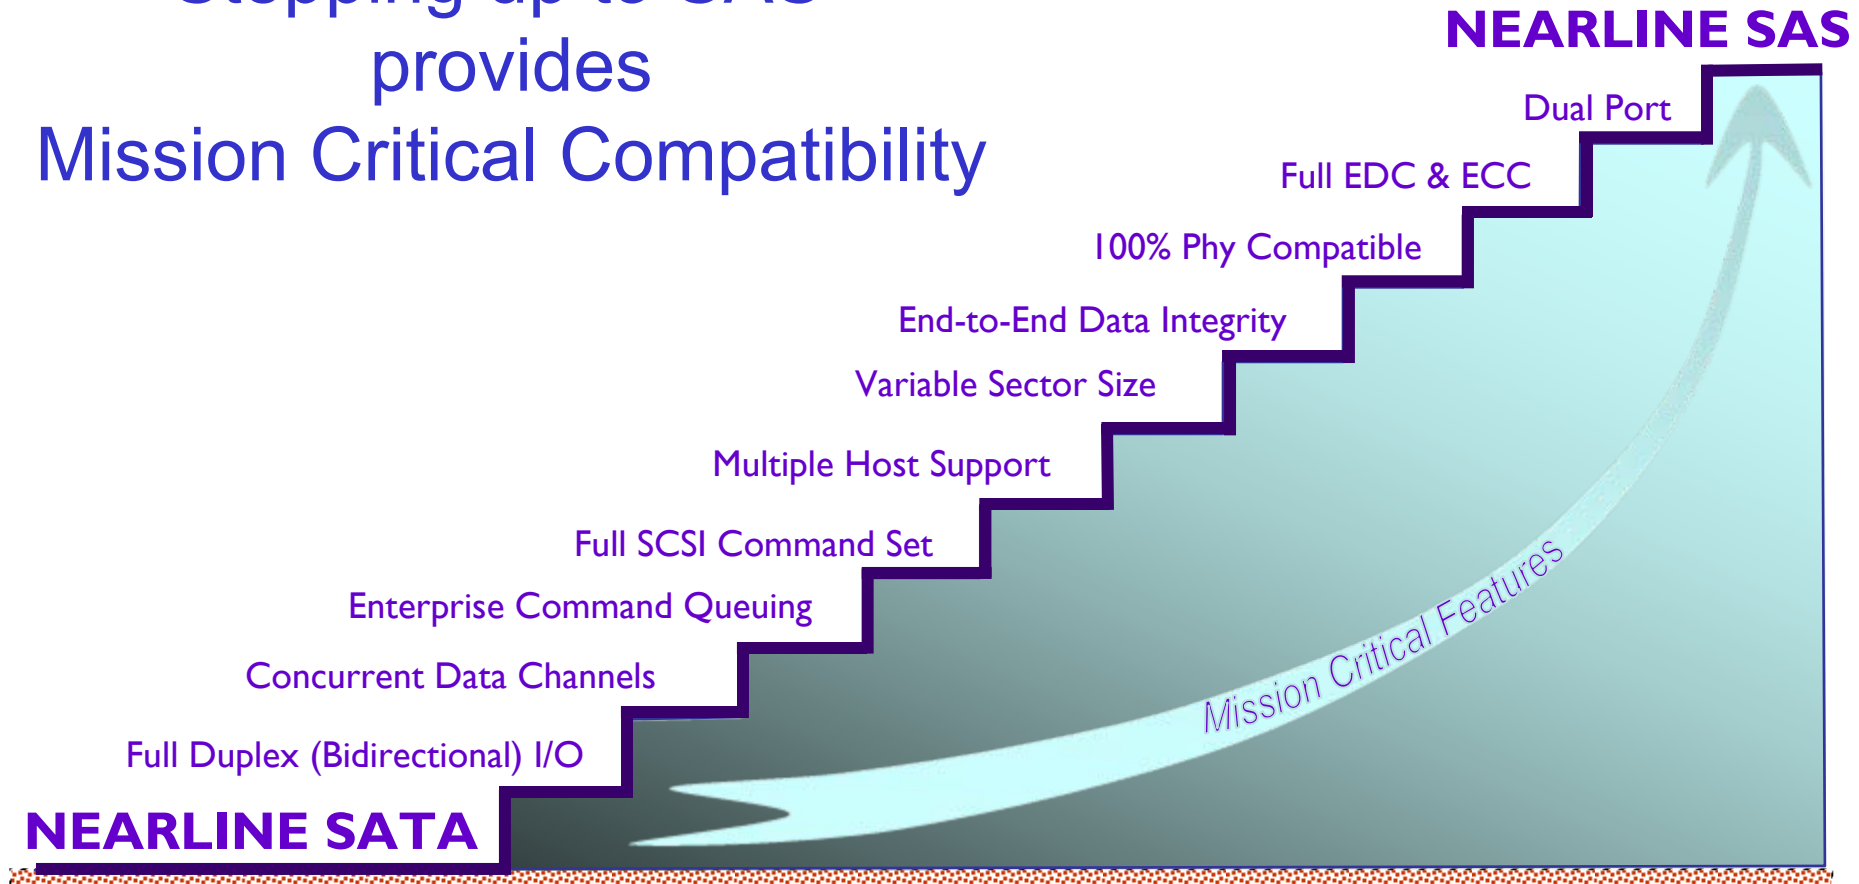


This prevents a SAS drive from being plugged into a SATA cabinet slot.

The Advantages of Nearline SAS

NL SATA Compared to NL SAS

Stepping up to SAS
provides
Mission Critical Compatibility



SAS/SATA NL Physical Differences




SAS
Electronics

SAS
Port "B"

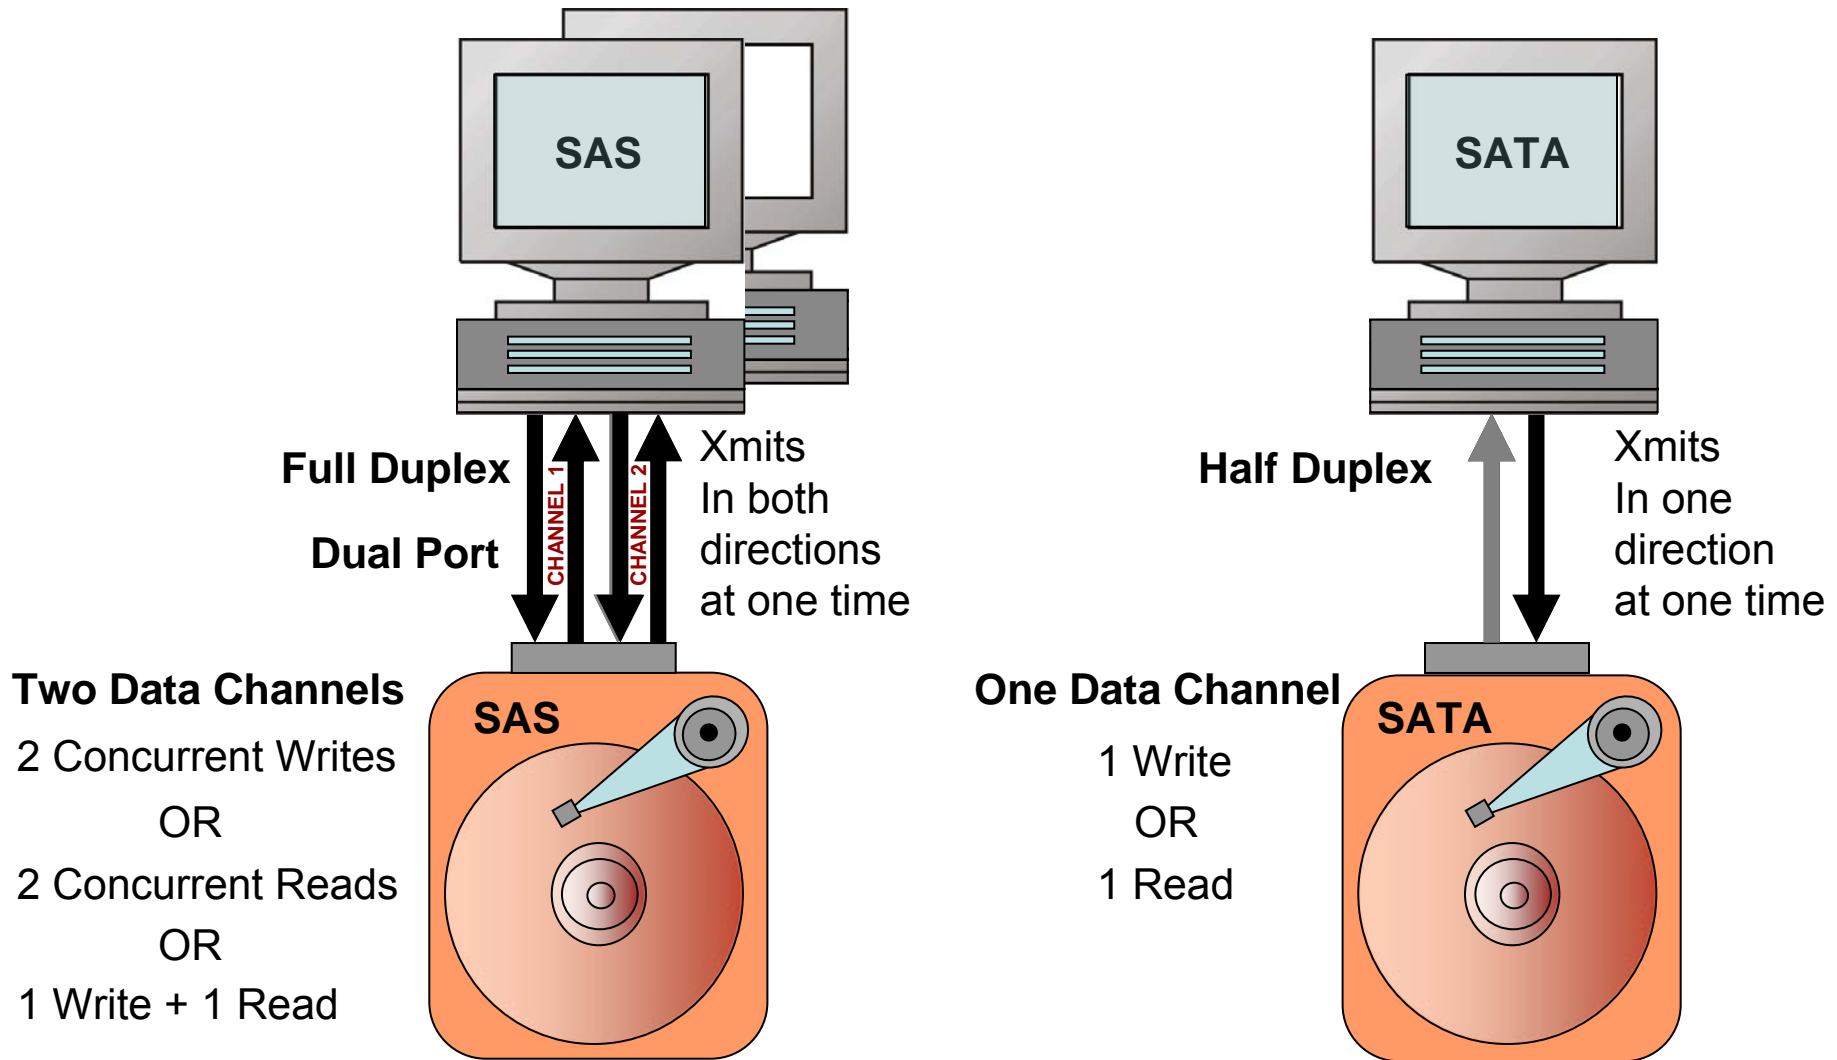


Nearline
Head/Disc Assy.





SATA
Electronics

SAS/SATA NL Differences



Multi Host Command Queuing





NL SAS/SATA Summary

- Both SATA and SAS Nearline drives are designed for use in Enterprise Mission Critical environments.
- SAS Nearline drives have additional advantages which are made possible by the Serial SCSI interface and enterprise electronics:
 - ◆ Full system interface compatibility at the protocol, physical (“phy”), and command level
 - ◆ Enterprise error recovery and performance optimization controls
 - ◆ Full data integrity protection both within the drive and at the system level with DIF¹ support.


¹Data Integrity Feature also known as T10 PI (Data Protection Information).

Rotational Vibration

RV Emitted by a Seeking Drive




HDA subjected to rotational forces




'Scope Picture, Seagate Prod. Dev.

Impact on Performance*



*Source: STX Competitive Analysis.

RV in 33 Different Cabinets




- More stringent RV spec. needed for SATA cabinets
- RV aggravated by system fans, random access and “bursty” workloads

Data Error Rate

UER* on High Capacity RAID Sets

- The UER for SATA desktop is 1 in 10^{14} bits transferred
 - ◆ 10^{14} bits = 12½ terabytes
- A 500 Gbyte drive has $1/25 \times 10^{14}$ bits
- Rebuilding a SATA drive in a RAID 5 set of 5 drives means transferring $5/25 \times 10^{14}$ bits = 1/5 of UER spec.
 - ◆ 20% probability of an Unrecoverable Error during the rebuild.
- Better odds would be available with RAID 1 or 6
 - ◆ RAID 1 rebuilds from a single mirror drive
 - ◆ RAID 6 can tolerate a second error during the rebuild.
- Risks can be reduced with good error management
 - ◆ Intelligent rebuild (ignore unused capacity)
 - ◆ Background media scan (dynamic certification)



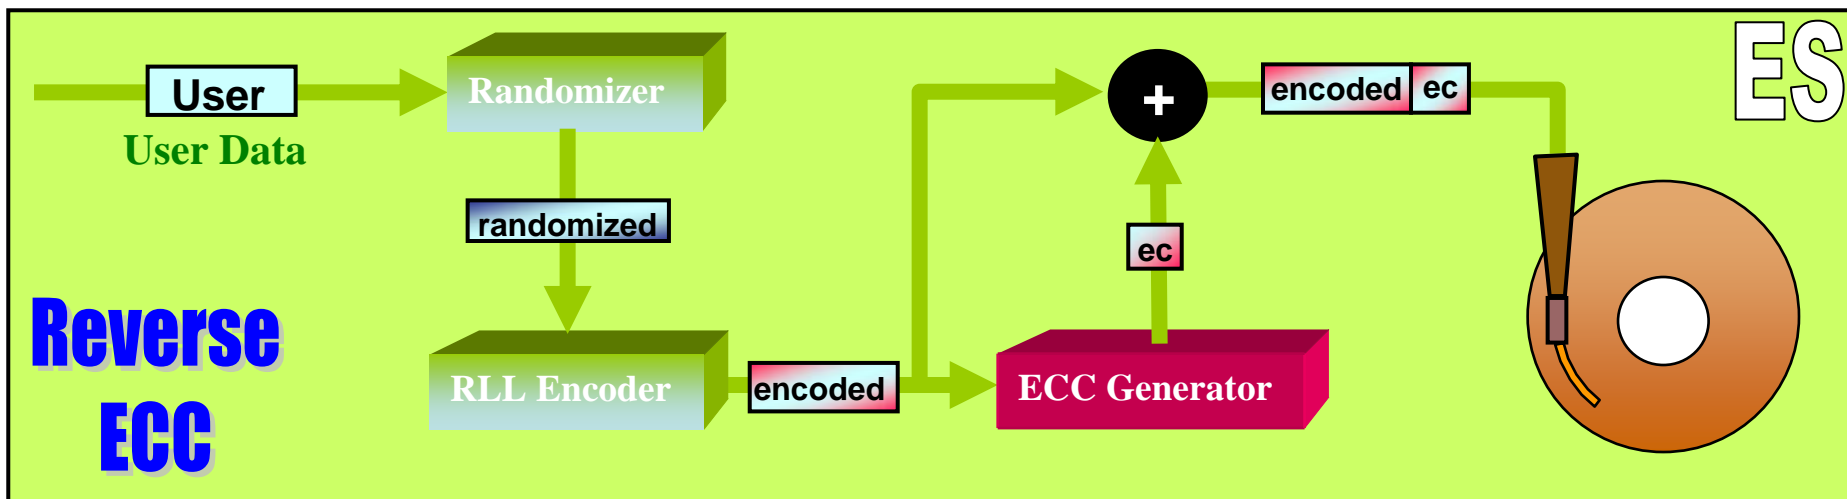
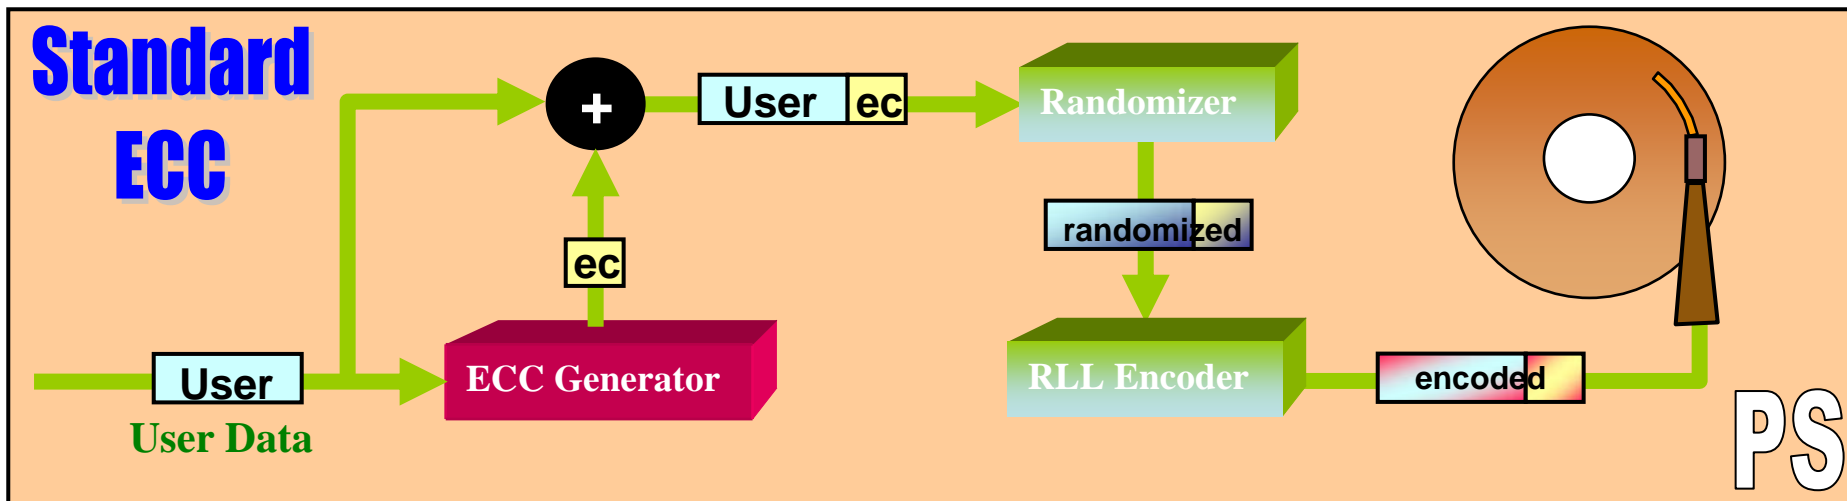
*Unrecoverable Error Rate

Desktop, Nearline & Enterprise HDDs – What's the difference?
© 2008 Storage Networking Industry Association. All Rights Reserved.

Error Correction Capability

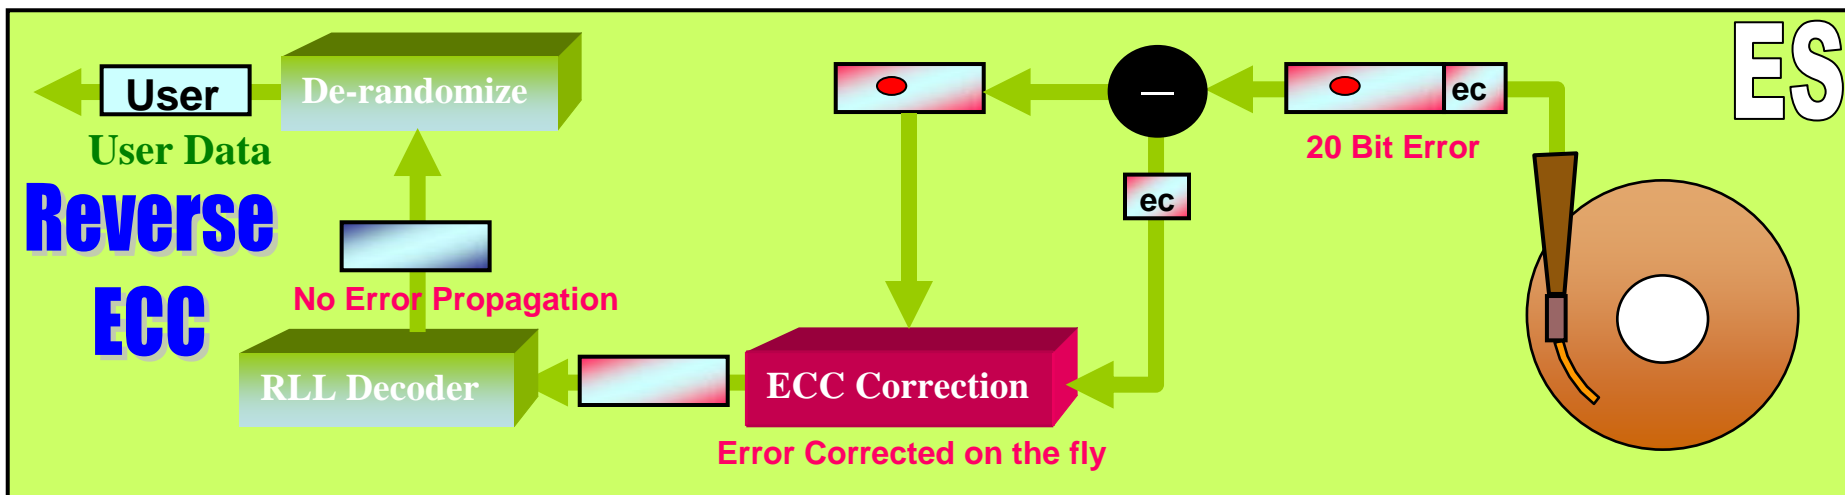
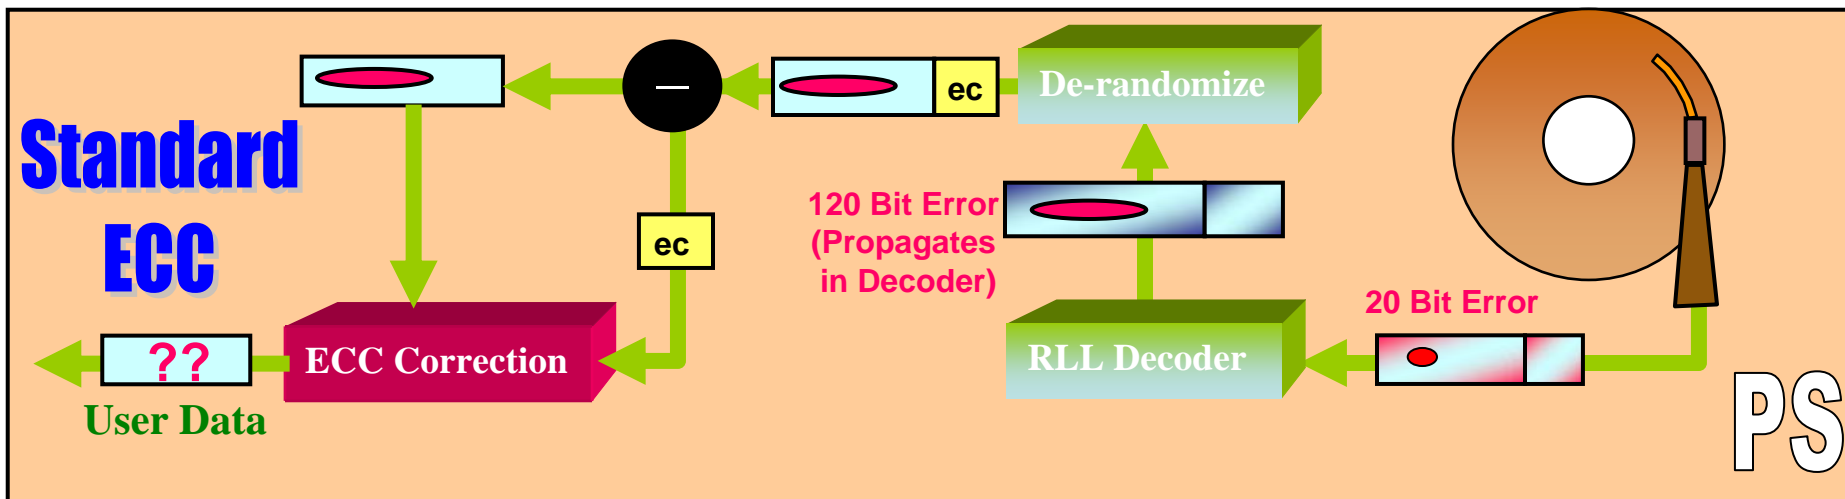
Standard vs Reverse ECC

(Write Command)



Standard vs Reverse ECC

(Read Command)

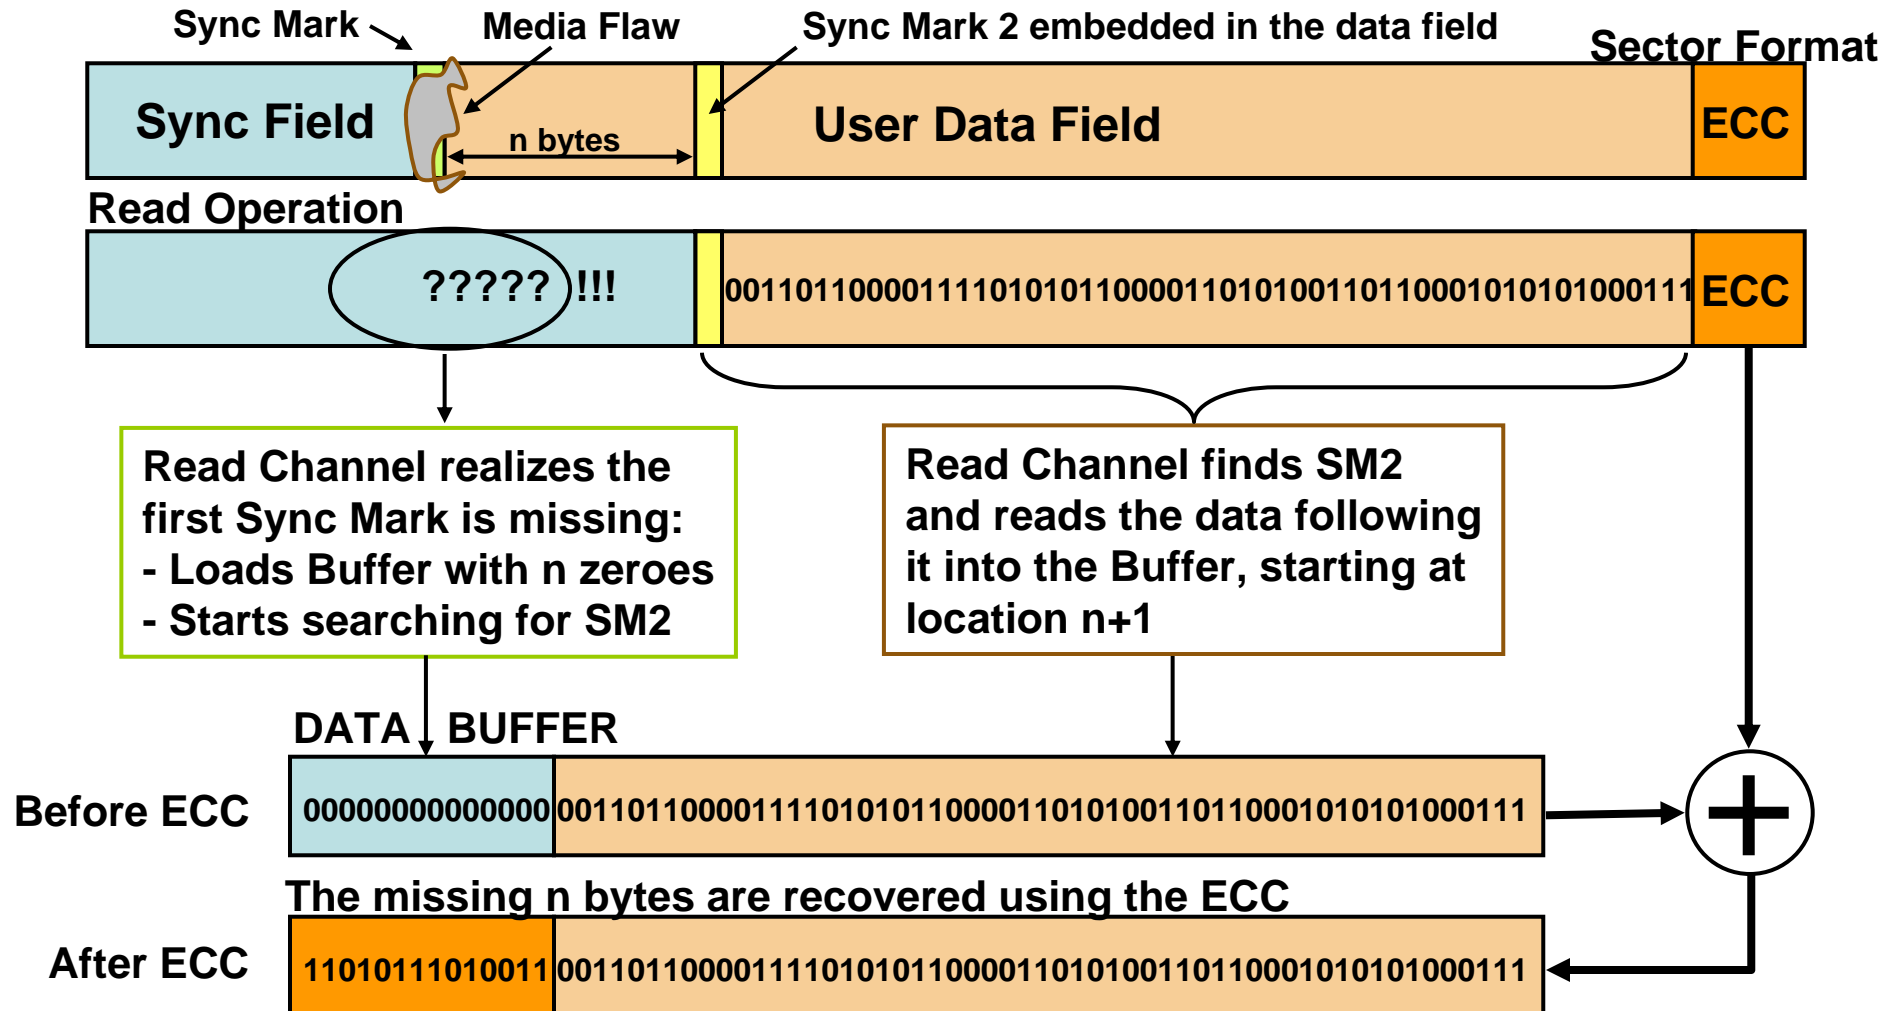


Sync Mark Errors on SATA Drives



- The Sync Field is used to get the read channel in frequency sync with the data recorded on the media
- The Sync Mark is used to define the beginning of the User Data Field
- Failure to recognize the Sync Mark (due to a thermal asperity or a grown media defect) means the User Data Field is not delineated and the data is lost.

Sync Mark Errors on SAS Drives



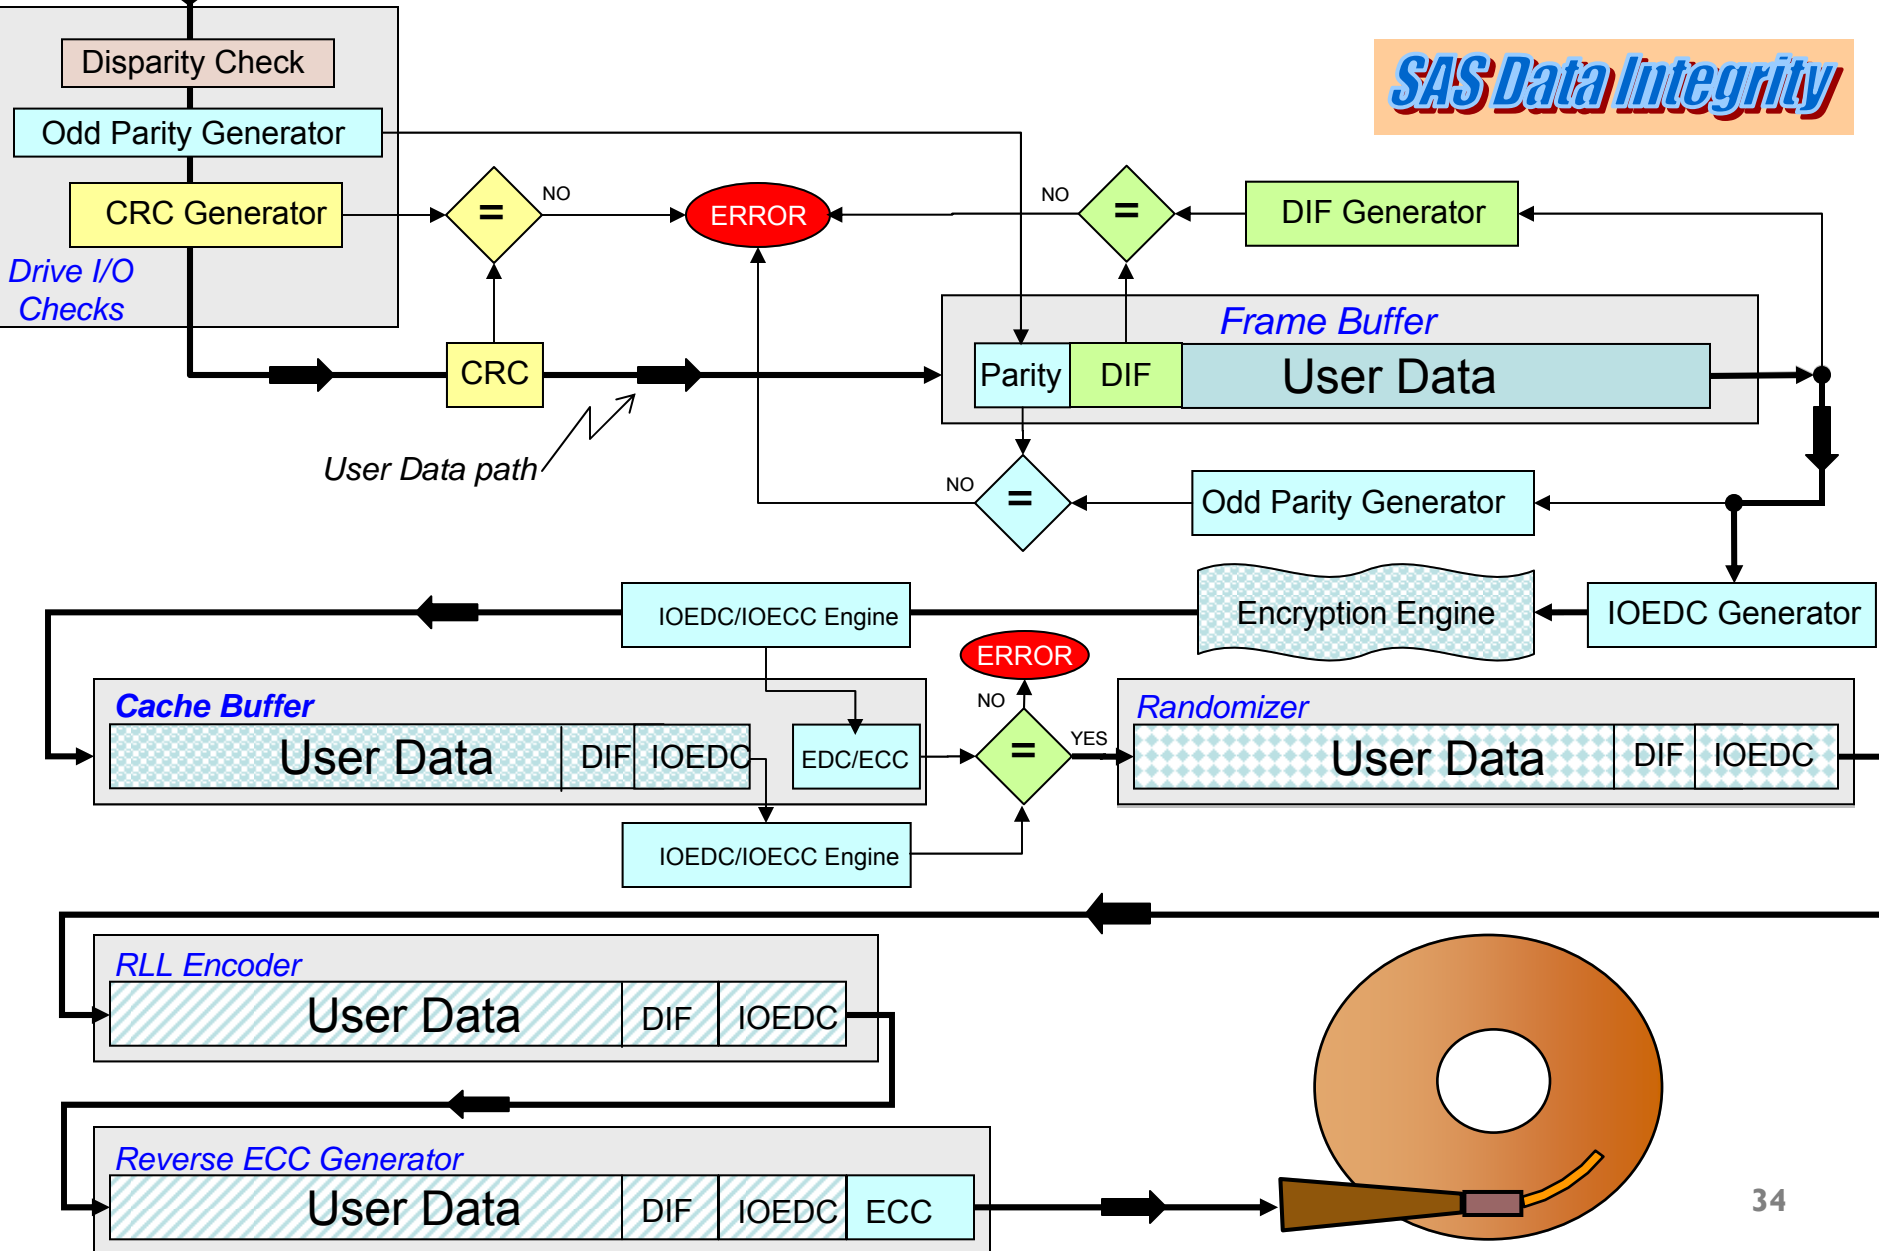
Data Integrity

Single Sector Data Transfer from the Host



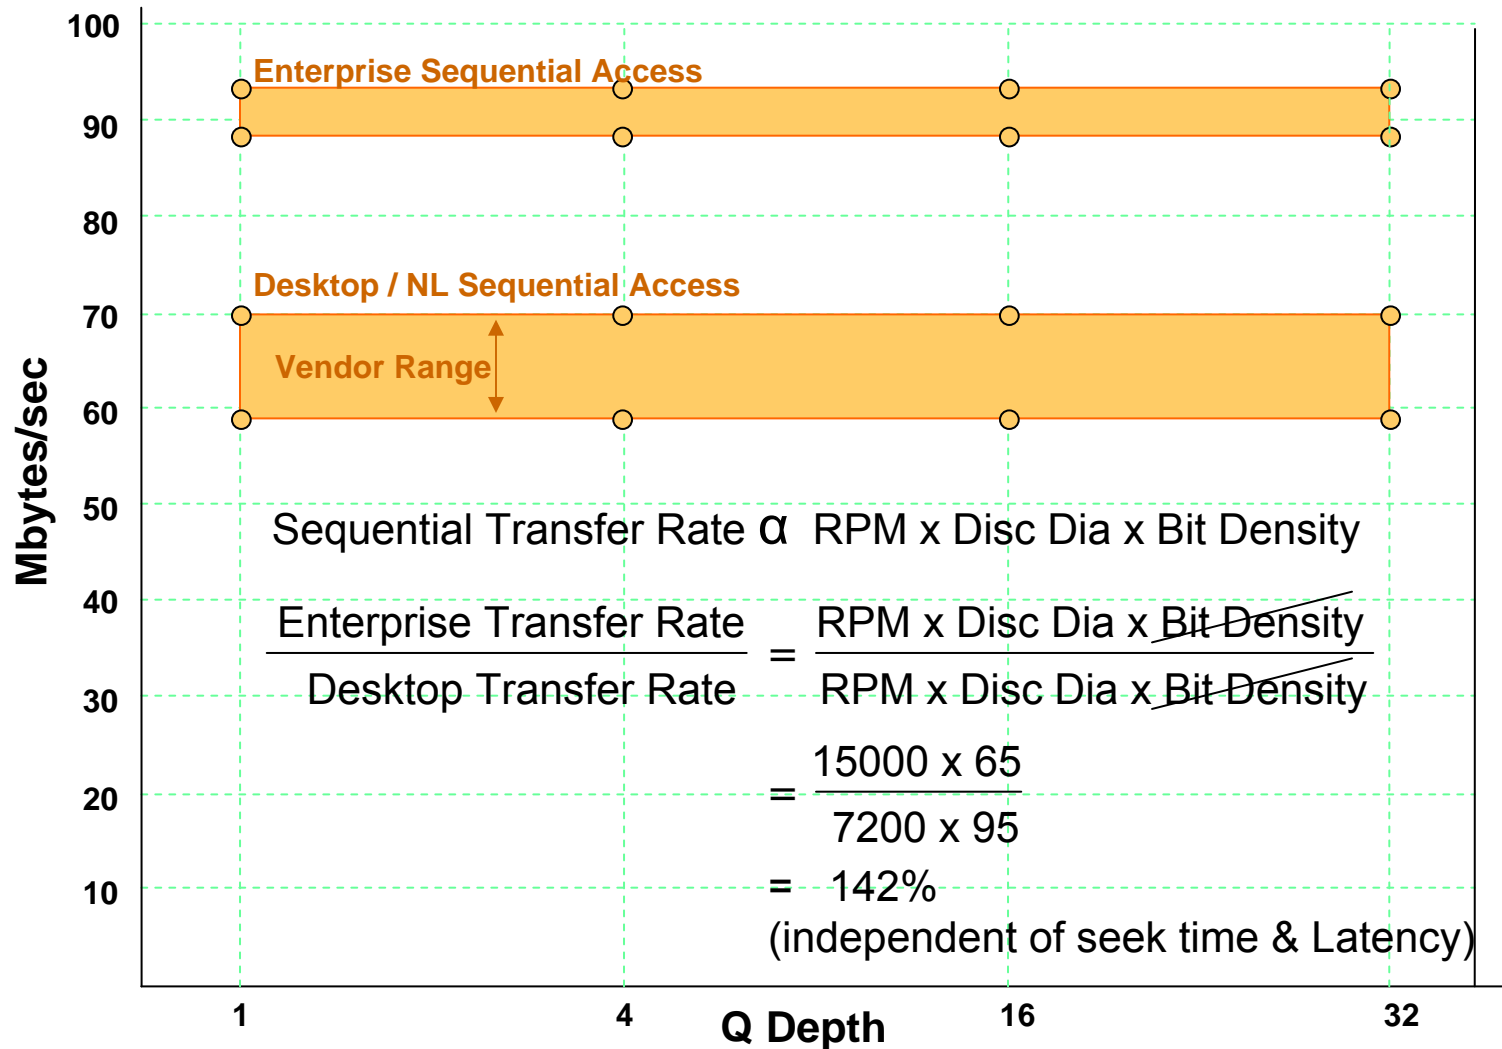
FC or SAS Interface

SAS Data Integrity




Performance


Performance Comparison




Performance Comparison



SPC-1C Performance Comparison




- SPC-1C comprises I/O operations designed to demonstrate small storage subsystem performance (1-16 drives) while performing the typical functions of a business critical application.



AFR

(Annualized Failure Rate)

AFR vs Duty Cycle



Although technological advances, driven by Enterprise research, will be leveraged into SATA products, there will continue to be functional limitations imposed on these devices by the overriding metric of **Low \$/GB Storage.**

- Please send any questions or comments on this presentation to SNIA:
trackstorage@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

SNIA Education Committee

**Craig Parris
Daniel Dummer
Willis Whittington
Wolfgang Rosner**